



US009313193B1

(12) **United States Patent**  
**Mehta et al.**

(10) **Patent No.:** **US 9,313,193 B1**  
(45) **Date of Patent:** **Apr. 12, 2016**

(54) **MANAGEMENT AND AUTHENTICATION IN HOSTED DIRECTORY SERVICE**

(71) Applicant: **Amazon Technologies, Inc.**, Seattle, WA (US)

(72) Inventors: **Gaurang Pankaj Mehta**, Seattle, WA (US); **Neelam Satish Agrawal**, Seattle, WA (US); **Lawrence Hun-Gi Aung**, Seattle, WA (US); **Guruprakash Bangalore Rao**, Bellevue, WA (US); **Shuo Wang**, Bellevue, WA (US); **Sameer Palande**, Seattle, WA (US); **Krithi Rai**, Redmond, WA (US); **Chirag Pravin Pandya**, Sammamish, WA (US)

(73) Assignee: **Amazon Technologies, Inc.**, Seattle, WA (US)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 5 days.

(21) Appl. No.: **14/500,865**

(22) Filed: **Sep. 29, 2014**

(51) **Int. Cl.**  
**H04L 29/06** (2006.01)  
**G06F 17/30** (2006.01)

(52) **U.S. Cl.**  
CPC ..... **H04L 63/08** (2013.01); **G06F 17/30386** (2013.01); **H04L 63/105** (2013.01)

(58) **Field of Classification Search**  
CPC ..... H04L 63/08; H04L 63/10  
See application file for complete search history.

(56) **References Cited**

**U.S. PATENT DOCUMENTS**

2009/0111428 A1\* 4/2009 Blommaert ..... H04W 12/04 455/411  
2010/0169952 A1\* 7/2010 Maki ..... H04L 63/0846 726/3

2011/0271099 A1\* 11/2011 Preiss ..... H04L 63/0807 713/155  
2012/0102553 A1\* 4/2012 Hsueh ..... H04L 9/3234 726/6  
2013/0054968 A1\* 2/2013 Gupta ..... H04L 9/3213 713/168  
2013/0097686 A1\* 4/2013 Towata ..... G06F 21/608 726/9  
2013/0252583 A1\* 9/2013 Brown ..... H04W 12/06 455/411  
2014/0230020 A1\* 8/2014 Mogaki ..... H04L 63/10 726/4  
2014/0237580 A1\* 8/2014 Kato ..... H04L 63/08 726/9  
2015/0135257 A1 5/2015 Shah et al.

**OTHER PUBLICATIONS**

Bindiganavale et al, Role Based Access Control in Enterprise Application, 2006, IEEE, pp. 111-116.\*

\* cited by examiner

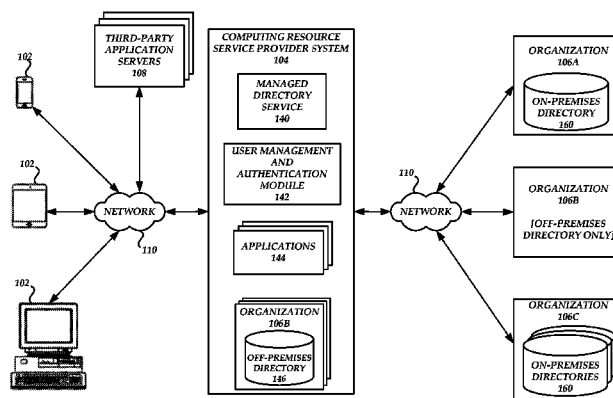
*Primary Examiner* — Izunna Okeke  
*Assistant Examiner* — Ali Shayanfar

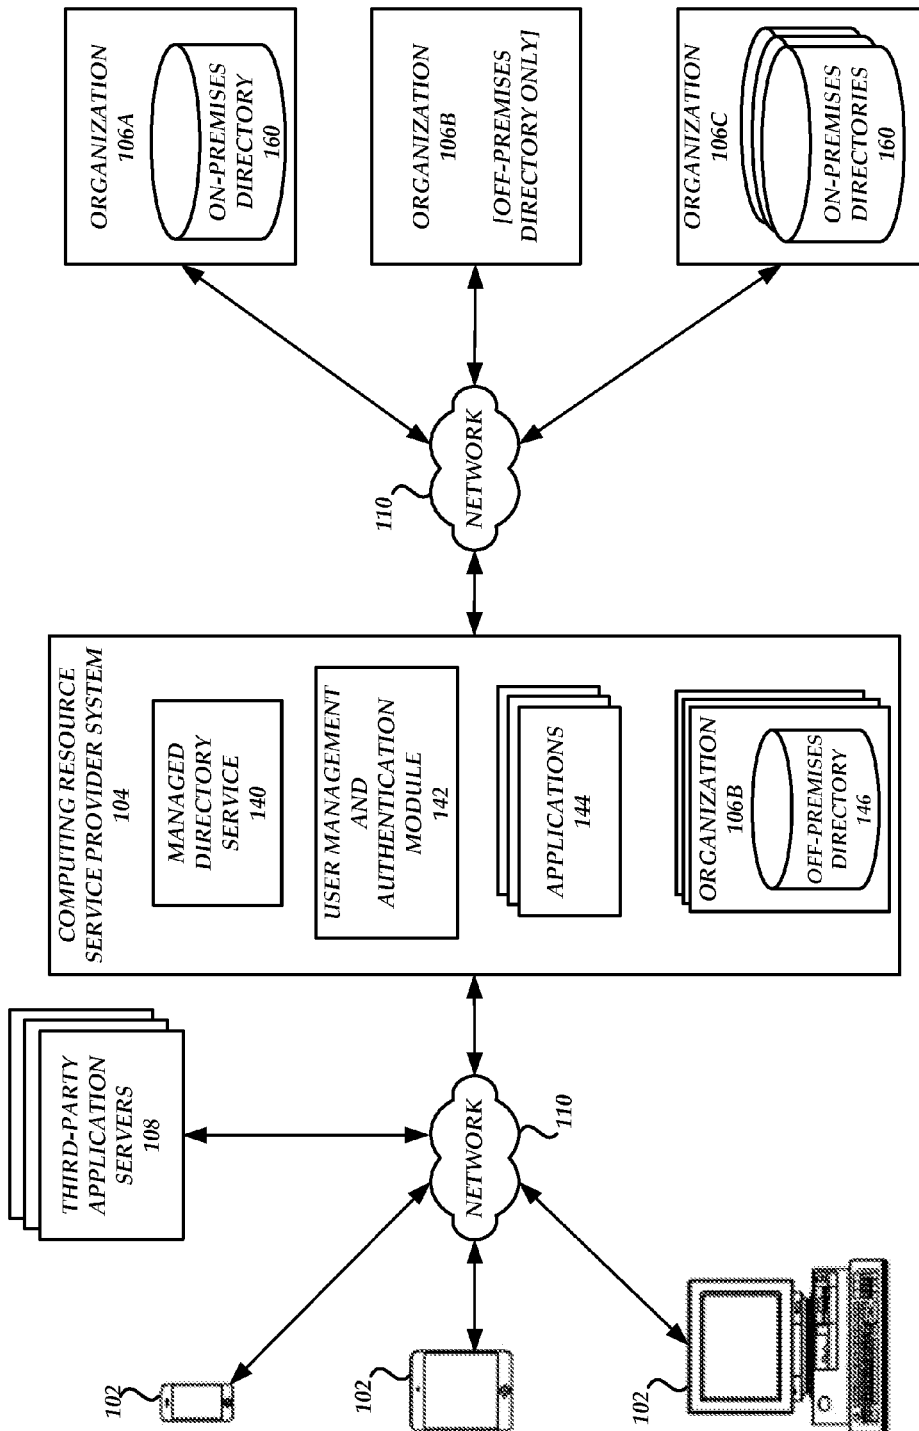
(74) *Attorney, Agent, or Firm* — Knobbe, Martens, Olson & Bear, LLP

(57) **ABSTRACT**

A user, group, and device management and authentication system allows administrators to manage one or more directories with devices that are not associated with a domain of the one or more directories via a set of APIs. The system also allows applications and services that do not have direct access to a list of directory users to access the one or more directories. The user, group, and device management and authentication system may be an add-on system that works in conjunction with a centrally-managed directory service to provide such functionality. For example, the system may generate an access token associated with a particular directory that can be used by a service accessed by an administrator to call an API provided by the system. The API call may be translated into a directory-specific API call that can be used to perform an action in the particular directory.

**13 Claims, 5 Drawing Sheets**





*Fig. 1*

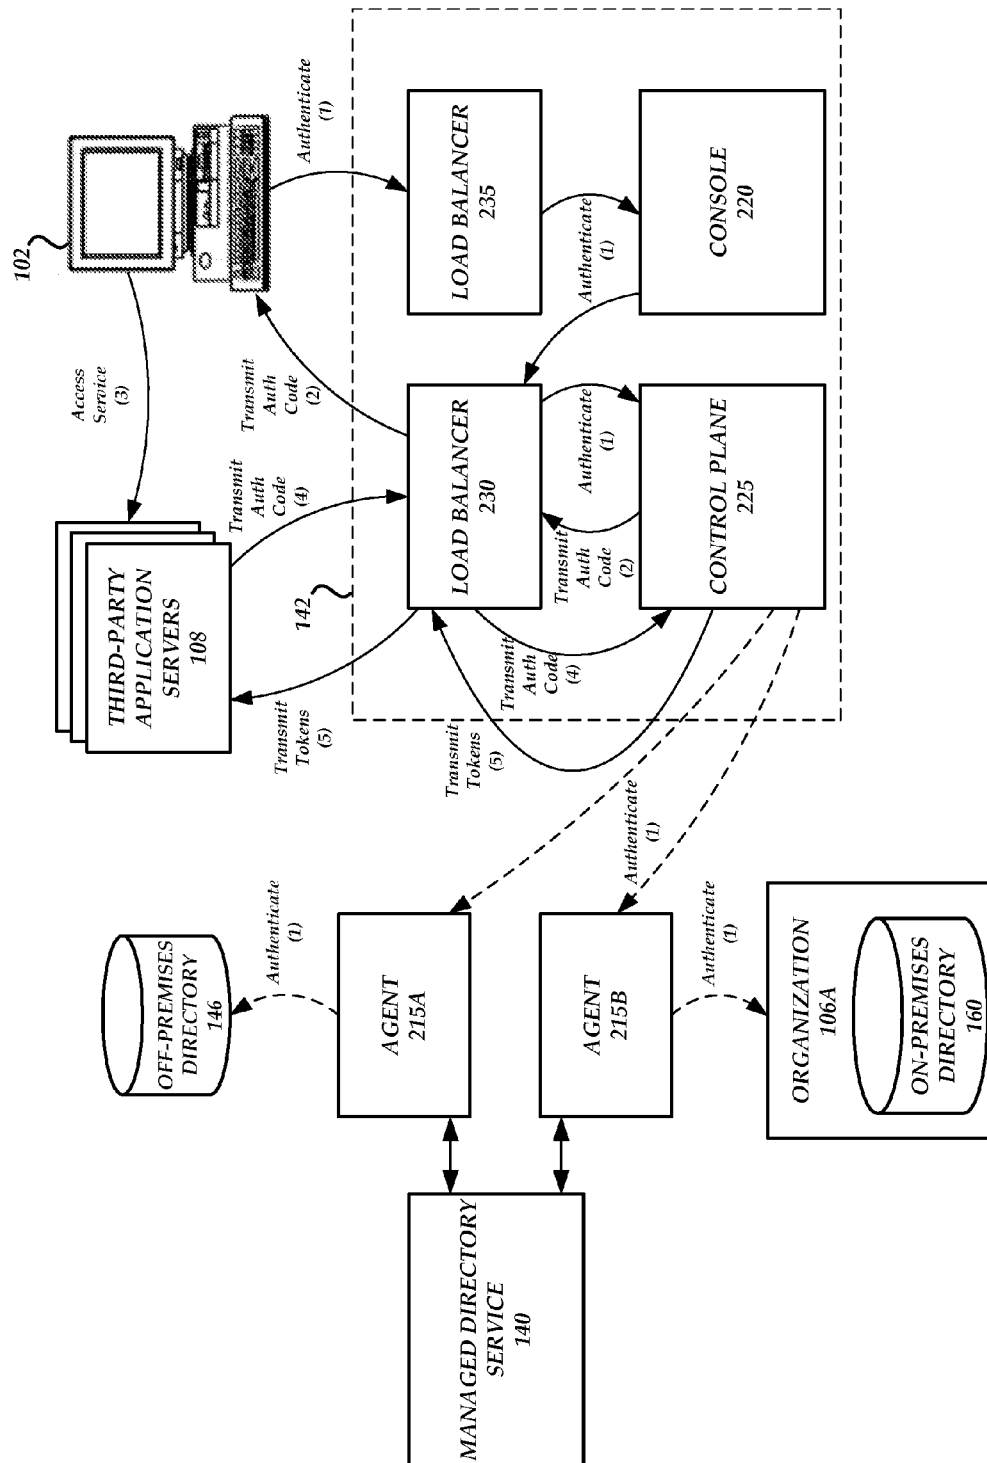


Fig. 2A

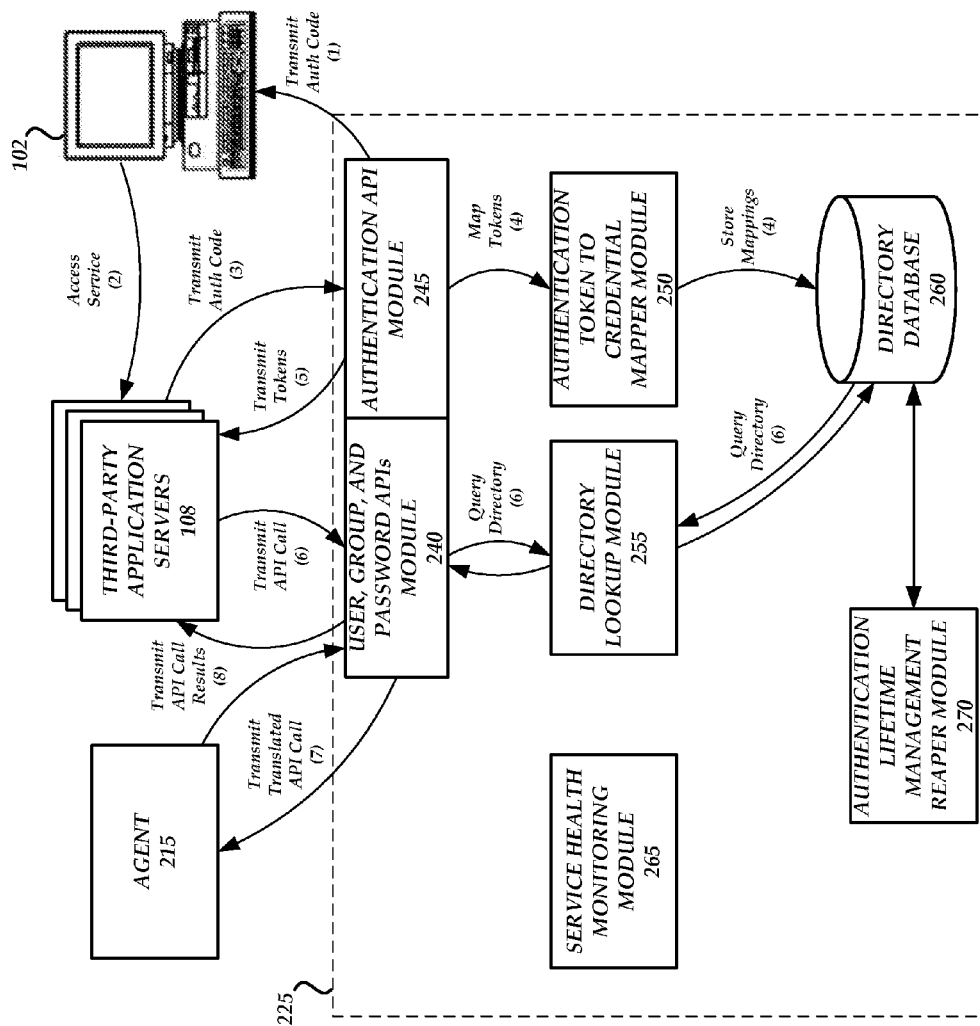
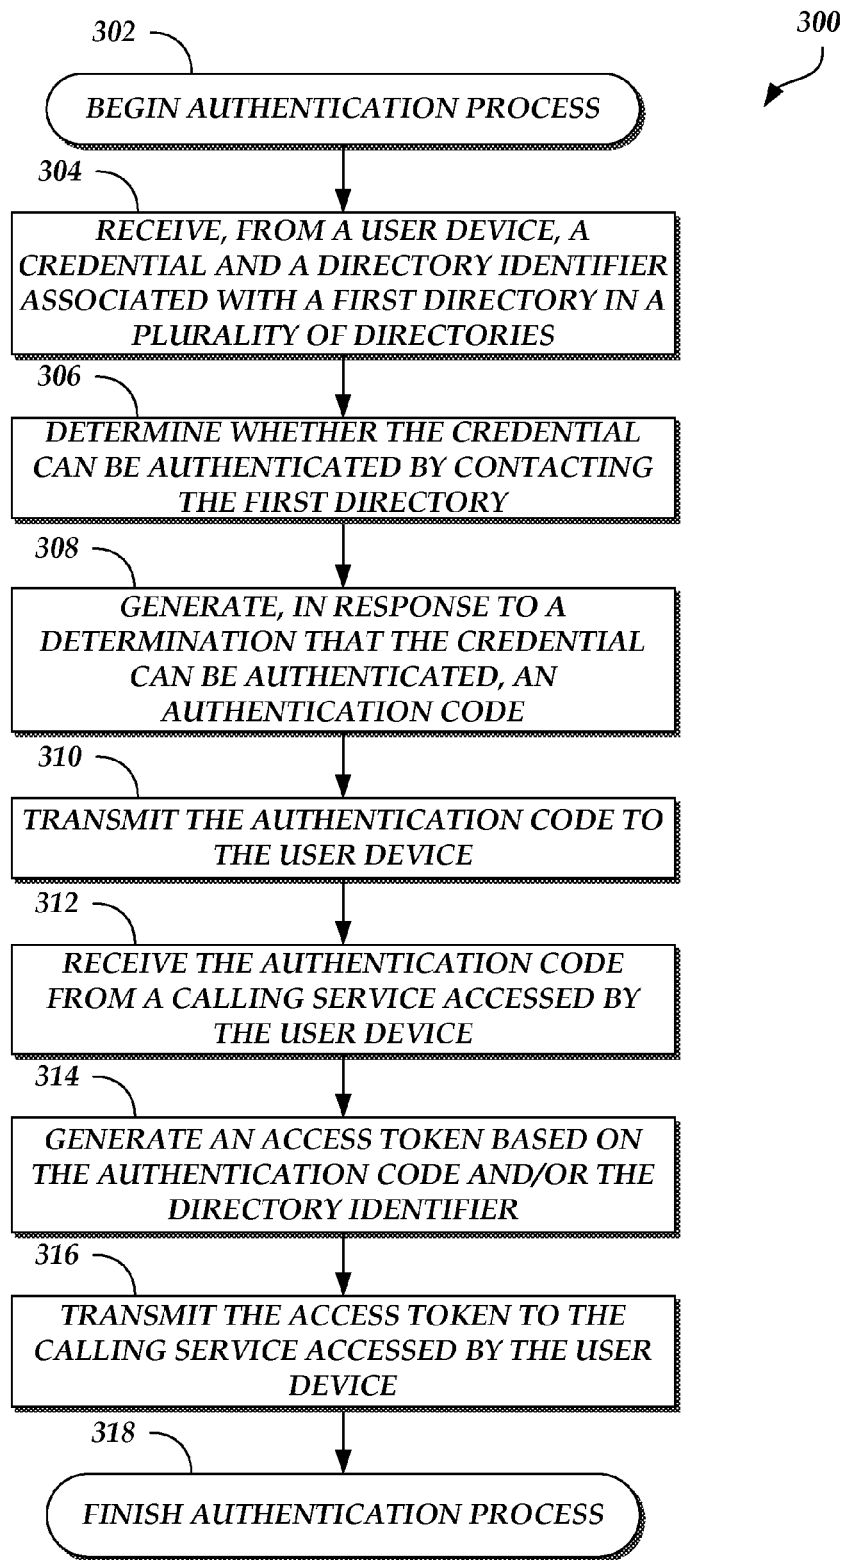
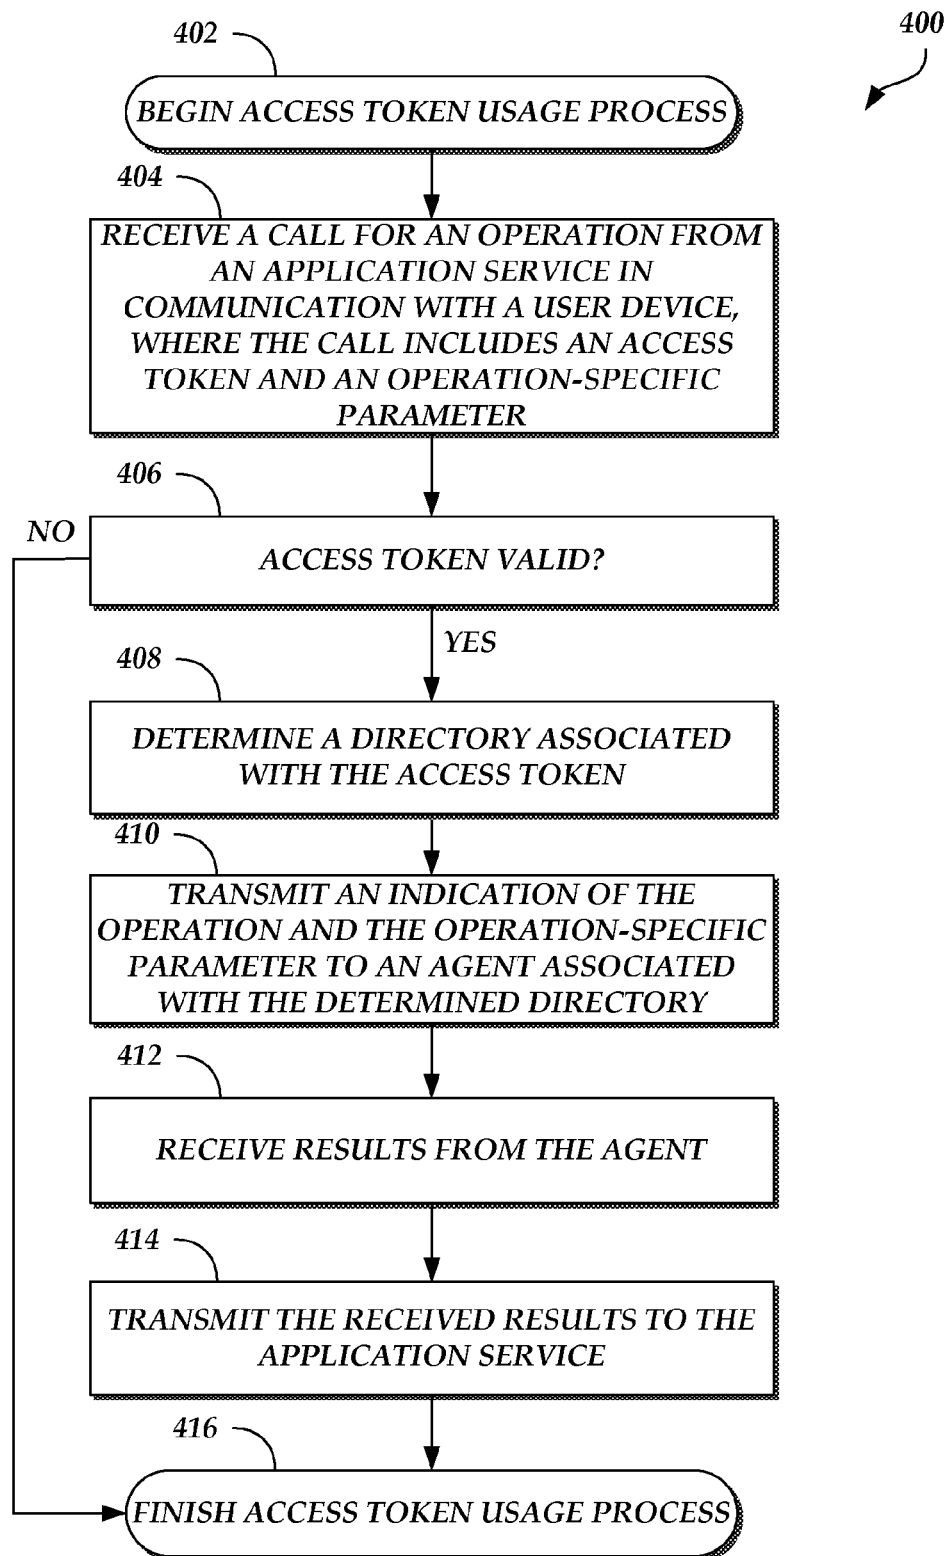


Fig. 2B

**Fig. 3**

**Fig. 4**

1

## MANAGEMENT AND AUTHENTICATION IN HOSTED DIRECTORY SERVICE

### BACKGROUND

Administrators may utilize directory services to create and maintain a directory for user, group, device, and/or computing resource management and/or for providing access to a variety of computing resources (e.g., file systems, files, users, security policies, network resources, applications, system storage, etc.). For example, the directory service may be implemented in a data server operated by an administrator (e.g., on-premises). The administrator may also operate a plurality of client devices, each of which shares a network or domain with the data server. To keep client devices secure and to ensure compatibility across the domain, the data server may assign and enforce security policies on the client devices and install or update software running on the client devices.

However, installing, maintaining, and operating the data server can be burdensome. The data server itself may include several computing systems, thus requiring the purchase of expensive hardware and the configuring of complex software. In some cases, dedicated facilities for powering and cooling the data server may be needed as well. Establishing and maintaining connectivity between the data server and the client devices may require the installation of expensive network equipment. Furthermore, additional hardware and/or software may be needed to implement backup and recovery procedures in case the data server fails or data is otherwise lost.

### BRIEF DESCRIPTION OF THE DRAWINGS

Throughout the drawings, reference numbers may be used to indicate correspondence between referenced elements. The drawings are provided to illustrate example embodiments described herein and are not intended to limit the scope of the disclosure.

FIG. 1 shows an example network environment in which directory management features and user, group, and device management and authentication features of the present disclosure can be implemented according to some embodiments.

FIG. 2A shows interactions to and from the user management and authentication module during the process of providing a service accessed by a user device with tokens.

FIG. 2B shows interactions to and from the user management and authentication module during the process of providing tokens and handling API calls.

FIG. 3 illustrates an authentication process that may be used by a computing resource service provider system to authenticate an administrator and provide the administrator with an access token that can be used to call APIs provided by the computing resource service provider system to manage one or more directories.

FIG. 4 illustrates an access token usage process that may be used by a computing resource service provider system to allow an administrator to manage one or more directories.

### DETAILED DESCRIPTION

#### Introduction

As described above, on-premises data servers that implement directory services can be burdensome. Thus, in some conventional systems, a managed directory service (e.g., a system that stores, organizes, and provides access to information in a computer operating system's directory, such as

2

MICROSOFT® ACTIVE DIRECTORY®) can be implemented by a centrally-managed data server that is located remotely and shared by a plurality of administrators and/or organizations. The managed directory service may manage a plurality of directories. The centrally-managed data server may provide access to the managed directory service via a network (e.g., the Internet) and an administrator may use existing network-enabled client devices to manage the directory. For example, the administrator may establish an account within a directory, and add member servers (e.g., a server that implements lightweight directory access protocol (LDAP), Kerberos, a domain name system (DNS) service, or other Active Directory tools to manage records in a directory) to the domain. Using a member server added to the domain, the administrator can create and manage domain users. Accordingly, the administrator can implement a directory service without having to install or maintain the infrastructure normally used to implement the directory service.

However, existing managed directory services have several limitations. For example, as described above, user management may be performed using a member server added to the domain. Some administrators, though, may not have access to a member server or may not wish to operate a member server (e.g., web developers, application developers, mobile application developers, etc.), and thus may not be able to perform any user management or otherwise manage the resources of a directory. As another example, an administrator may have created multiple directories, but the managed directory service may only allow the administrator to access and manage one directory at a time. Applications or services (whether running on a user's device or accessible via a network) that do not have direct access to the list of created users in a domain or directory may not be able to access the domain or directory. Thus, it may not be possible for an administrator or user using an application or service associated with one domain or directory to share resources or content with another domain (e.g., a domain that does not share the same tree or forest as the domain associated with the application or service) or directory.

Accordingly, the embodiments described herein present a user management and authentication system that (1) allows administrators to perform user management with devices that are not associated with a domain and to access a plurality of directories or domains via a set of application programming interfaces (APIs), and (2) provides authorization and authentication mechanisms for allowing applications and services that do not have direct access to a list of created users to access a domain or directory. The user management and authentication system may be an add-on system that works in conjunction with the managed directory service to provide the functionality described herein. For example, the user management and authentication system may generate a set of login pages (e.g., content or network pages, such as web pages) that can be accessed by a device associated with an administrator. The administrator can enter his or her credentials, a client identification that identifies an application or service that is requesting access to a domain, a redirect page (e.g., a page associated with an application or service that the user management and authentication system should instruct the user device to access once login is complete), and/or a directory identifier (or domain identifier). The user management and authentication

tion system may transmit the credentials to the directory associated with the directory identifier, and the directory may determine whether the credentials can be authenticated (e.g., whether the administrator has access to the directory). If the credentials can be authenticated, the user management and authentication system is notified and generates an authentication code (e.g., an OAuth code).

The authentication code may be transmitted by the user management and authentication system to the user device along with an instruction to access the redirect page. The authentication code may be a single-use code that is valid for a set period of time (e.g., 10 minutes, 1 hour, etc.) and, before expiration of the authentication code, may be used by the application or service associated with the redirect page to initiate access to the directory associated with the directory identifier. For example, the user management and authentication system may include a getToken API. The application or service may call the getToken API, passing the authentication code as a parameter.

The getToken API may generate an access token and/or a refresh token in response to receiving a valid authentication code and provide the tokens to the application or service. The access token and/or the refresh token may be generated based on the credentials and/or the directory identifier. For example, the access token and/or the refresh token may include the credentials and/or the directory identifier in a secure format. The access token may allow the application or service to access and/or manage the directory indicated by the directory identifier via a set of APIs and may be valid for an administrator-defined or preset period of time (e.g., 1 hour, 1 day, etc.). The refresh token may be valid for an administrator-defined or preset period of time (e.g., 1 week, 1 month, etc.) and can be used by the application or service to receive a new access token once the previous access token expires. In some embodiments, the refresh token may be valid for no amount of time (e.g., the access token may not be refreshed once it expires).

In an embodiment, the user management and authentication system provides several APIs. Such APIs may include user APIs, group APIs, organizational unit APIs, password APIs, access token APIs, and/or service APIs. User APIs may include a createUser API (e.g., to create a user in a directory), a describeUsers API (e.g., to list all or any number of users within a directory and their attributes), an updateUser API (e.g., to update the attributes of a user in a directory), a deleteUser API (e.g., to delete a user from a directory), and/or a listGroupsForUser API (e.g., to list all or any number of users within a directory and their groups). Group APIs may include a createGroup API (e.g., to create a new group within a directory), a describeGroups API (e.g., to list groups and attributes of the groups within a directory), an updateGroup API (e.g., to update an existing group in a directory), a deleteGroup API (e.g., to delete a group from a directory), a listMembersInGroup API (e.g., to list the members of groups in the directory), an addMemberToGroup API (e.g., to add members (users or groups) to a group in a directory), and/or a removeMemberFromGroup API (e.g., to remove a member from a group in a directory). Organization unit APIs may include a describeOrganizationalUnits API (e.g., to list all or any number of organizational units within a directory and their attributes). Password APIs may include an authenticateUser API (e.g., to authenticate a user in the directory and return an authentication code), an authenticateKerberosUser API (e.g., to authenticate a user in the directory and return an authentication code), an authenticateRadiusUser API (e.g., to authenticate a user against a radius server associated with a directory and return an authentication code), a resetPassword

API (e.g., to reset a user's password), and/or a changePassword API (e.g., to change a user's password). Access Token APIs may include a createAnonymousToken API (e.g., to create an anonymous token to store), a getToken API (e.g., to generate an access and/or refresh token based on an authentication code), a validateToken API (e.g., to verify that a previously issued access token or anonymous token is still valid), a refreshToken API (e.g., to generate a new access token using a previously issued refresh token), and/or a revokeToken API (e.g., to invalidate a previously issued access, anonymous, or refresh token). Service APIs may include a getServiceAccountCreds API (e.g., to allow a registered application or service to retrieve domain join credentials).

The application or service can use the access token to call any of the APIs supported by the user management and authentication system. For example, the access token and/or other operation-specific parameters can be provided to the user management and authentication system via an API to manage a directory. As described above, the access token may be generated based on the authentication code and/or the directory identifier. Thus, when unwrapped by the user management and authentication system upon reception via the API, the access token may identify the authentication code, and thus the administrator that is performing the action, and/or the directory that the action is to be performed on. Accordingly, the application or service may not separately identify the directory to be accessed and/or managed. Not separately identifying the directory may provide an additional level of security not found in conventional systems. In conventional systems, which generally involve accessing and managing a single directory, an administrator may manage the directory simply by providing a directory identifier and/or other operation-specific parameters. However, in the embodiments disclosed herein, a user or administrator may not be able to access and manipulate a directory in an unauthorized manner simply by providing a directory identifier and/or other operation-specific parameters. Rather, the access token is needed to access and/or manipulate a specific directory, and the user or administrator may obtain the access token only after his or her credentials can be verified with the specific directory (e.g., after the authentication code is received), as described herein.

In an embodiment, when generating the access token, the user management and authentication system may map the access token to a directory service token (e.g., a Kerberos token, other directory-specific credentials, etc.) that normally could be used by a member server to access the directory identified by the directory identifier. Thus, when an application or service calls an API using the access token, the user management and authentication system may access a database to determine the token that maps to the access token, and pass the mapped token (and/or other operation-specific parameters provided by the application or service) to the managed directory service to perform the action or operation requested via the calling of the API. In some embodiments, the validity of the refresh token may be tied to the validity of the mapped token (e.g., the lifetime of which can be preset or defined by an administrator). Results, if any, may be returned by the managed directory service to the user management and authentication system, and the user management and authentication system may forward the results to the application or service.

The application or service can repeat this process to generate access tokens for one or more directories operated by the administrator. Thus, the administrator, using an application or service, can access and/or manage a plurality of directories.



5

In some embodiments, a directory service and a user management and authentication system is assigned to a region (e.g., a physical location). A plurality of directory services and user management and authentication systems may be present, separated by regions. The directory service in one region may not interact with or access a directory service or user management and authentication system in another region. Likewise, the user management and authentication system in one region may not interact with or access a directory service or user management and authentication system in another region. In other embodiments, a directory service and a user management and authentication system is assigned to a region, but the directory service and/or the user management and authentication system may interact with directory services and/or user management and authentication systems in other regions.

In further embodiments, the user management and authentication system supports single-factor authentication and/or multi-factor authentication. The user management and authentication system may support either type of authentication via user interfaces and/or APIs. The user management and authentication system may also support single sign-on (e.g., an administrator may be able to login once and gain access to all appropriate directories without being prompted to log in again each time a directory is accessed).

Example Network Environment

FIG. 1 shows an example network environment in which directory management features and user, group, and device management and authentication features of the present disclosure can be implemented according to some embodiments. As used herein the term “directory” generally refers to an organized collection of data about users, devices, applications, and other common resources of a computer network. Each resource on a computer network (or some subset thereof) may be represented as an object in a directory, and information about a particular resource (e.g., name, address, permissions, etc.) can be stored as attributes of that object. Information can be securely stored within or in association with the object such that only users with sufficient permissions are able to access, modify, or otherwise use the information.

As shown, the network environment includes various user devices **102**, a computing resource service provider system **104**, organizations **106** and third-party application servers **108** in communication via one or more networks **110**. The computing resource service provider system **104** can provide applications; directory management services; user, group, and device management and authentication services; and/or other network-based services to various organizations or other customers. Organizations **106A-C** (or other customers) can employ the computing resource service provider system **104** to provide application access to users associated with the organizations, manage the organizations’ directories, etc. Individual users can use user devices **102** to access applications hosted by the computing resource service provider system **104** (or third-party application servers **108**) using credentials from their respective organizations **106A-106C**. In addition, the computing resource service provider system **104** can provide the applications with access to the directories of the various organizations **106A-C** at the discretion of the respective organizations.

The user devices **102** can correspond to a wide variety of computing devices, including desktop computing devices, laptop computing devices, terminal devices, mobile phones, tablet computing devices, media players, wearable computing devices (e.g., smart watches, smart eyewear, etc.), and various other electronic computing devices and appliances

6

having one or more computer processors, computer-readable memory and network-access capabilities. Some user devices **102** may be associated with a particular organization **106A-C**. For example, an organization may have various user devices **102** that remain on-premises, or that are used off-premises primarily by employees, administrators, or other users associated with the organization. In some embodiments, some or all of the user devices **102** may be separate from any organization, such as public computers or home computers that are used by any number of users to perform various tasks, which may include managing directories or accessing applications using credentials associated with a particular organization **106A-C** or other customer of the computing resource service provider system **104**.

Individual user devices **102** may execute an application to communicate via the network **110** with the computing resource service provider system **104** in order to manage one or more directories and/or to access provided applications or services. For example, the application may be a stand-alone application that is installed on the user device **102**. As another example, the application may be a browser (e.g., a web browser) that accesses an application or service (e.g., a web service) hosted by the computing resource service provider system **104**, the third-party application server **108**, and/or another computing system (not shown).

The computing resource service provider system **104** can be a computing system configured to host or otherwise provide access to applications **144** (word processing applications, photo-editing applications, electronic mail applications, etc.), manage directories for separate customer organizations **106A-C**, and/or provide other network-based services and resources (e.g., document-sharing services, virtual machine services, etc.). For example, the computing resource service provider system **104** can be a server or group of servers that may be accessed via a communication network **110**. The computing resource service provider system **104** can include a number of components to provide various features described herein, such as a managed directory system or service **140**, a user management and authentication module **142**, and one or more applications or application servers **144** that can be accessed by organizations **106** and user devices **102**. The computing resource service provider system **104** may also store various off-premises directories **146**, such as an off-premises directory for organization **160B**, as described below. In some embodiments, the computing resource service provider system **104** may include additional or fewer components than illustrated in FIG. 1 to provide the features described above and in greater detail below.

As used herein, the term “off-premises directory” refers to a directory that is remote from the organization with which it is associated, in order to distinguish such a directory from a directory that is located on an organization’s premises. Thus, although a directory may be physically stored on the premises of a computing resource service provider system **104**, the directory may nevertheless be referred to as an off-premises directory because it is off-premises with respect to the organization with to which it belongs (e.g., the organization that owns or operates the network described by the directory). Additionally, although a directory may be physically stored off the premises of the computing resource service provider system **104**, the directory may nevertheless be referred to as an on-premises directory because it is on-premises with respect to the organization to which it belongs.

Illustratively, an administrator may use the application executed by the user device **102** to manage one or more directories owned or operated by the administrator’s organization, such as one of organizations **106A-C**. The application

may interact with the managed directory service **140** and/or the user management and authentication module **142**. The managed directory service **140** may be a computing system that implements a managed directory service. In an embodiment, the managed directory service **140** is configured to create, monitor, and manage one or more directories. For example, the managed directory service **140** may be in communication with and manage the off-premises directory **146** and/or the on-premises directories **160**. As described above, an administrator may use the managed directory service **140** to create, monitor, and/or manage a directory if the user device **102** is a member server. However, if the user device **102** is not a member server, the administrator may create, monitor, and/or manage the directory via APIs provided by the user management and authentication module **142**.

The user management and authentication module **142** may be a computing system that implements a user, group, and device management and authentication system. In an embodiment, the user management and authentication module **142** allows administrators to manage one or more directories with user devices **102** that are not member servers (e.g., that are not associated with a domain of the respective directory) via a set of APIs, such as the APIs described above. The user management and authentication module **142** may also provide authorization and authentication mechanisms for allowing the application executed by the user device **102** or applications or services accessed by the executed application to access content or resources in a directory even if the executed application or the accessed applications or services do not have direct access to a list of created users of the directory. For example, the user management and authentication module **142** may be in communication with the managed directory service **140** and may serve as an interface between the user device **102** and the managed directory service **140** such that the user device **102** can manage one or more directories managed by the managed directory service **140**. The user device **102** can call an API provided by the user management and authentication module **142**, and the user management and authentication module **142** can instruct the managed directory service **140** to perform an action indicated by the called API. The interaction between the user device **102**, the managed directory service **140**, and the user management and authentication module **142** is described in greater detail below with respect to FIGS. 2A-2B.

In further embodiments, the user management and authentication module **142** supports single-factor authentication and/or multi-factor authentication. The user management and authentication module **142** may support either type of authentication via user interfaces and/or APIs. The user management and authentication module **142** may also support single sign-on (e.g., an administrator may be able to login once and gain access to all appropriate directories without being prompted to log in again each time a directory is accessed).

In an embodiment, applications or services provided by the third-party application servers **108** or the computing resource service provider system **104** that have registered with the computing resource service provider system **104** may access the features described herein. Applications or services that have not registered with the computing resource service provider system **104** and/or that have not been approved by administrators of the computing resource service provider system **104** may be barred from accessing the features described herein.

The computing resource service provider system **104** may be a single computing device, or it may include multiple distinct computing devices, such as computer servers, logically or physically grouped together to collectively operate as

a server system. The components of the computing resource service provider system **104** can each be implemented in application-specific hardware (e.g., a server computing device with one or more ASICs) such that no software is necessary, or as a combination of hardware and software. In addition, the modules and components of the computing resource service provider system **104** can be combined on one server computing device or separated individually or into groups on several server computing devices.

In addition, multiple (e.g., two or more) computing resource service provider systems **104** may be used. For example, computing resource service provider systems **104** may be located in separate regions and may or may not interact with each other. The separate computing resource service provider systems **104** may be located so that they are close (in either a geographical or networking sense) to groups of current or potential user devices **102** or organizations **160A-C**.

In some embodiments, the features and services provided by the computing resource service provider system **104** may be implemented as web services consumable via the communication network **110**. In further embodiments, the computing resource service provider system **104** is provided by one more virtual machines implemented in a hosted computing environment. The hosted computing environment may include one or more rapidly provisioned and released computing resources, which computing resources may include computing, networking and/or storage devices. A hosted computing environment may also be referred to as a cloud computing environment.

The organizations **106A-C** can correspond to various customers of the computing resource service provider system **104**. Although the term "organization" is used herein, the features involving such organizations may additionally or alternatively involve any customer having a directory (whether on-premises or off-premises) and wishing to use the computing resource service provider system **104** to manage the directory and control access to the directory by applications hosted by the computing resource service provider **104** or third-party application servers **108**.

Organizations that maintain on-premises directories **160** may have one or more servers on which the directories **160** are stored. For example, organization **106A** may have a data center that includes various servers, and an on-premises directory **160** may be stored on one or more of the servers. Organizations that maintain off-premises directories may employ the services of the computing resource service provider system **104**, which may store the off-premises directory in an off-premises directory data store **146**. For example, organization **106B** may not maintain an on-premises directory at all, but may rely instead on the computing resource service provider system **104** to maintain the organization's directory **146**. Some organizations may choose to maintain multiple directories on-premises and/or off-premises. For example, organization **106C** may store multiple on-premises directories **160**, each in a manner similar to organization **106A** (described above), and the organization **106C** may also choose to employ the computing resource service provider system **104** to maintain an off-premises directory **146**. The directory **146** maintained by the computing resource service provider system **104** in this example may be a mirror or subset of the on-premises directory (e.g. for backup or disaster-recovery purposes), or it may be a separate directory altogether (e.g., a directory of computing resources in a different region from the on-premises directory **160**).

The communication network **110** may be a publicly accessible network of linked networks, possibly operated by vari-

ous distinct parties, such as the Internet. In some embodiments, the communication network **110** may be or include the Internet, a private network, personal area network, local area network, wide area network, cable network, satellite network, cellular telephone network, etc. or combination thereof.

Example Interactions to and from the User Management and Authentication Module

FIG. 2A shows interactions to and from the user management and authentication module **142** during the process of providing a service accessed by a user device with tokens. As illustrated in FIG. 2A, the directory service module **140** may interact with a plurality of agents **215A-B**, and the agents **215A-B** may interact with the user management and authentication module **142**.

In an embodiment, each agent **215A-B** is associated with one or more separate directories and directly interfaces with its associated directories for management purposes. The agents **215A-B** may be associated with and communicate with on-premises and/or off-premises directories. For example, the agent **215A** may be associated with the off-premises directory **146** and the agent **215B** may be associated with the on-premises directory **160**. The managed directory service **140** may be configured to create, monitor, and/or manage the agents **215A-B**. While two agents **215A-B** are illustrated, this is not meant to be limiting. The computing resource service provider system **104** may include any number of agents (e.g., a number of agents sufficient to handle all directories managed by the managed directory service **140**).

The agents **215A-B** may receive translated versions of API calls made to the user management and authentication module **142**. The translated versions of the API calls made to the user management and authentication module **142** may be directory-specific API calls (e.g., LDAP, Kerberos, DNS, etc.) that can be executed by the managed directory service **140**. As an example, translation of the API calls may include mapping an access token to a directory service token (e.g., a Kerberos token, a username and password pair, an NT LAN manager (NTLM) hash, etc.).

In an embodiment, the user management and authentication module **142** includes a console **220**, a control plane **225**, and load balancers **230** and **235**. The console **220** may be configured to generate user interfaces that are transmitted to the user devices **102**. The user interfaces may be login pages that may be transmitted to a user device **102** when a user device **102** calls an API provided by the user management and authentication module **142** to login, reset a password, change a password, and/or perform other operations described herein. The console **220** may also generate a link (e.g., a uniform resource locator (URL)) that can be transmitted to the user device **102**, such as, for example, when a new user is created. The link may be valid for a finite period of time (e.g., 7 days, 2 weeks, etc.) and, when selected, may redirect the application executed by the user device **102** to a content page that allows the administrator to enter additional information, such as user profile information. The console **220** may also transmit an electronic message (e.g., email, text message, etc.) that includes the link to an account associated with an administrator when a user device **102** calls an API provided by the user management and authentication module **142** to reset a password. The electronic message may include a one-time user token that can be used by the user device **102** to complete the password reset process.

The control plane **225** may be configured to expose APIs to the user devices **102**. For example, the control plane **225** may expose APIs like the APIs described herein. The control plane **225** may also be configured to interact with the agents **215A-B**. For example, the control plane **225** may translate the API

calls received from the user devices **102** into directory-specific API calls that can be executed by the managed directory service **140** and provide the directory-specific API calls to the appropriate agent **215A-B** (e.g., the agent associated with the directory on which an action is to be performed according to the API call). The control plane **225** is described in greater detail below with respect to FIG. 2B.

The console **220** and the control plane **225** may each be behind a load balancer **230** or **235**. The console **220** and the control plane **225** may each include multiple computing resources and the load balancers **230** and **235** may distribute workloads across the multiple computing resources to optimize resource use, to maximize throughput, and/or to minimize the risk that any single resources becomes overloaded. For example, the load balancers **230** and **235** may receive API calls from the user device **102** and distribute the API calls to the appropriate computing resources of the console **220** or the control plane **225**.

At (1), an administrator, via the user device **102**, may first communicate with the load balancer **235** to authenticate his or her credentials. The load balancer **235** may forward the authentication request and credentials to the console **220**, which may forward the authentication request and credentials to the load balancer **230**. The load balancer **230** may forward the authentication request and credentials to the control plane **225**. The control plane **225** may determine a directory associated with the credentials and transmit the authentication request and credentials to the appropriate agent **215A** or **215B**. Once the appropriate agent **215A** or **215B** receives the authentication request and credentials, the agent **215A** or **215B** may perform the authentication by contacting its associated directory (e.g., off-premises directory **146** or on-premises directory **160**).

At (2), in response to a determination that the administrator's credentials can be authenticated, the control plane **225** generates an authentication code. The control plane **225** transmits the authentication code to the load balancer **230** for forwarding to the user device **102**.

At (3), the user device **102** accesses a service provided by the third-party application servers **108**. In alternative embodiments, not shown, the user device **102** accesses an application in the applications **144** provided by the computer resource service provider system **104**.

At (4), the accessed service transmits the authentication code to the load balancer **230** for the purpose of receiving an access token and/or a refresh token. The load balancer **230** may forward the authentication code to the control plane **225**.

At (5), the control plane **225** generates the access token and/or the refresh token based on the received authentication code. The control plane **225** may generate the access token if the authentication code is received before expiration. The control plane **225** may transmit the access token and/or the refresh token to the load balancer **230** for forwarding to the accessed service.

FIG. 2B shows interactions to and from the user management and authentication module **142** during the process of providing tokens and handling API calls. As illustrated in FIG. 2B, the control plane **225** may include a user, group, and password APIs module **240**, an authentication API module **245**, an authentication token to credential mapper module **250**, a directory lookup module **255**, a directory database **260**, a service health monitoring module **265**, and an authentication lifetime management reaper module **270**. In alternate embodiments, not shown, the directory lookup module **255** is a component of an agent **215A** or **215B**.

The authentication API module **245** may generate an authentication code, a refresh token, and/or an access token.

## 11

For example, the authentication API module **245** may receive administrator credentials, a client identification that identifies an application or service hosted by the applications server **140** that is requesting access to a domain, a redirect page (e.g., a page associated with an application or service that the user management and authentication module **150** should instruct the user device **102** to access once authentication is complete), and/or a directory identifier (or domain identifier). The authentication API module **245** may transmit the credentials to the agent **215A** or **215B** that corresponds with a directory associated with the directory identifier. The agent **215A** or **215B** may pass the credentials to the directory service module **140** to determine whether the credentials can be authenticated (e.g., whether the administrator has access to the directory). If the credentials can be authenticated, the authentication API module **245** is notified and generates an authentication code (e.g., an OAuth code).

At (1), as described above, the authentication code may be a single-use code that is valid for a finite period of time (e.g., 10 minutes, 1 hour, etc.). The authentication API module **245** may transmit the authentication code to the user device **102**. The authentication API module **245** may also transmit to the user device **102** an instruction to access the redirect page.

At (2), the user device **102** accesses a service provided by the third-party application servers **108**. In alternative embodiments, not shown, the user device **102** accesses an application in the applications **144** provided by the computer resource service provider system **104**.

At (3), the authentication API module **245** receives the authentication code from the accessed service (e.g., accessed via a browser). For example, the authentication API module **245** may receive the authentication code if the getToken API is called (e.g., the authentication code may be included as parameter).

If an unexpired authentication code is received, the authentication API module **245** may generate an access token and/or a refresh token. The access token and/or the refresh token may be generated based on the credentials and/or the directory identifier. For example, when unwrapped, the access token and/or the refresh token may indicate the authentication code (and thus the credentials and administrator associated with the credentials) and the directory identifier associated with the token. The access token may allow the accessed service to manage the directory indicated by the directory identifier via a set of APIs provided by the user, group, and password APIs module **240**. The access token may be valid for an administrator-defined or finite period of time (e.g., 1 hour, 1 day, etc.). The refresh token may be valid for an administrator-defined or finite period of time (e.g., 1 week, 1 month, etc.) and can be used by the user device **102** and/or the accessed application or service to receive a new access token once the previous access token expires. In some embodiments, the refresh token may not be valid for any period of time (e.g., the access token may not be refreshed once it expires).

At (4), in some embodiments, the authentication API module **245** transmits the access token, the refresh token, the credentials, and/or the directory identifier to the authentication token to credential mapper module **250**. The authentication token to credential mapper module **250** may use underlying directory logic to map the credentials and/or the directory identifier to the access token and/or the refresh token and store this mapping in the directory database **260**.

At (5), the authentication API module **245** may also transmit the access token and/or the refresh token to the accessed service. The user, group, and password APIs module **240** may provide one or more of the APIs described above, and the accessed service may use the access token and/or other opera-

## 12

tion-specific parameters to call one or more of the provided APIs. In some embodiments, not shown, the authentication API module **245** transmits the access and/or refresh tokens to the accessed service before performing the operations discussed with respect to (4).

At (6), in an embodiment, upon receiving an API call that includes an access token from a calling service, the user, group, and password APIs module **240** identifies a directory associated with the access token by querying the directory lookup module **255**. For example, the directory lookup module **255** may pass the access token to the directory database **260** and identify a directory associated with the access token, transmitting the identified directory back to the user, group, and password APIs module **240**.

At (7), the user, group, and password APIs module **240** may then identify an agent **215A** or **215B** associated with the identified directory, translate the API call into a directory-specific API call, and pass the directory-specific API call and its operation-specific parameters to the appropriate agent **215A** or **215B**. The user, group, and password APIs module **240** may determine how to translate the API call into a directory-specific API call based on information provided by the directory lookup module **255**. In alternate embodiments, not shown, the user, group, and password APIs module **240** directly identifies a directory associated with the access token without querying the directory lookup module **255** (e.g., via an internal lookup table or database).

As described above, the agents **215A-B** may receive directory-specific API calls translated from API calls received by the user, group, and password APIs module **240**. As an example, the user, group, and password APIs module **240** may map the access token received by the user, group, and password APIs module **240** to a directory service token (e.g., a Kerberos token), and the directory service token may be provided to the agents **215A-B** to access the appropriate directory.

At (8), results, if any, may be received by the agents **215A-B** from the managed directory service **140** and forwarded to the user, group, and password APIs module **240**. The user, group, and password APIs module **240** may then transmit the results to the calling service.

The service health monitoring module **265** may be a computing system that monitors the health of a directory service. For example, the service health monitoring module **265** may monitor the health of the directory service module **140**.

The authentication lifetime management reaper module **270** may perform maintenance on the directory database **260**. For example, the authentication lifetime management reaper module **270** may reap out or remove access tokens and/or refresh tokens that have expired (and their associated credentials and directory identification).

Example Process for Authenticating an Administrator and Providing an Access Token

FIG. 3 illustrates an authentication process **300** that may be used by a computing resource service provider system to authenticate an administrator and provide the administrator with an access token that can be used to call APIs provided by the computing resource service provider system to manage one or more directories. As an example, the computing resource service provider system **104** (e.g., the user management and authentication module **142**) of FIG. 1 can be configured to execute the authentication process **300**. The authentication process **300** begins at block **302**.

At block **304**, a credential and a directory identifier associated with a first directory in a plurality of directories are received from a user device. In an embodiment, a client iden-

13

tification that identifies a service that is not associated with the first directory is also received.

At block **306**, whether the credential can be authenticated is determined. In an embodiment, the credential is authenticated by passing the credential to the first directory.

At block **308**, in response to a determination that the credential can be authenticated, an authentication code is generated. The authentication code may be an OAuth code and may be valid for a finite period of time.

At block **310**, the authentication code is transmitted to the user device. For example, the authentication code is received by an application (e.g., browser) executed by the user device. In other embodiments, the authentication code is transmitted to an application or service hosted by an applications server and accessed by the user device.

At block **312**, the authentication code is received from a calling service accessed by the user device. The calling service be an application or service that the user device has accessed (e.g., via the browser). The authentication code may be received before the authentication code expires.

At block **314**, an access token is generated based on the authentication code and/or the directory identifier. For example, when unwrapped, the access token may identify the authentication code (e.g., which may include the credential that identifies the administrator) and the directory identifier. The access token may be valid for a finite period of time. In some embodiments, a refresh token is also generated. The refresh token may be valid for a finite period of time and may be used to generate a new access token when the previous access token expires.

At block **316**, the access token is transmitted to the calling service accessed by the user device. In an embodiment, the user device uses the access token to allow an administrator to manage the directory corresponding to the directory identifier (e.g., the first directory) via the calling service. After the access token is transmitted to the calling service, the authentication process **300** may be complete, as shown in block **318**. Example Process for Using an Access Token

FIG. 4 illustrates an access token usage process **400** that may be used by a computing resource service provider system to allow an administrator to manage one or more directories. As an example, the computing resource service provider system **104** (e.g., the user management and authentication module **142**) of FIG. 1 can be configured to execute the access token usage process **400**. The access token usage process **400** begins at block **402**.

At block **404**, a call for an operation from an application service in communication with a user device is received, where the call includes an access token and an operation-specific parameter. For example, the user device, via an application or service, may call an API, such as the createUser API, passing the access token and a username as operation-specific parameters. The operation-specific parameters passed via the API may include all operation-specific parameters required by the API schema associated with the called operation.

At block **406**, the access token usage process **400** determines whether the access token is valid. For example, the access token usage process **400** determines whether the access token is valid in time (e.g., determines whether the access token has expired) and, if the access token is valid in time, whether the access token is valid for the application or the service that the user device is attempting to access (e.g., the managed directory service). If the access token is valid, the access token usage process **400** proceeds to block **408**. Otherwise, the access token usage process **400** finishes at block **416**.

14

At block **408**, a directory associated with the access token may be determined. In an embodiment, the access token is mapped to a directory when the access token is generated and the mapping is stored in a database, such as the directory database **260**. A directory lookup module, such as the directory lookup module **255** may query the directory database **260** to determine a directory associated with the received access token.

At block **410**, an indication of the operation and the operation-specific parameter are transmitted to an agent associated with the determined directory. In further embodiments, the user management and authentication module **142** maps the access token to a directory service token (e.g., a Kerberos token, a username and password pair, an NTLM hash, etc.) and the directory service token is also provided to the agent. In an embodiment, an agent is associated with a directory. The directory service token, an indication of the operation, and/or the operation-specific parameter may then be transmitted to the appropriate agent. The agent may use the received data to instruct the managed directory service, such as the managed directory service **140**, to perform the operation specified by the called API.

At block **412**, results from the agent are received. The results may be data received from the agent once the operation specified by the API call is complete.

At block **414**, the received results are transmitted to the application service. After the received results are transmitted to the application service, the access token usage process **400** may be complete, as shown in block **416**.

#### Terminology

All of the methods and tasks described herein may be performed and fully automated by a computer system. The computer system may, in some cases, include multiple distinct computers or computing devices (e.g., physical servers, workstations, storage arrays, cloud computing resources, etc.) that communicate and interoperate over a network to perform the described functions. Each such computing device typically includes a processor (or multiple processors) that executes program instructions or modules stored in a memory or other non-transitory computer-readable storage medium or device (e.g., solid state storage devices, disk drives, etc.). The various functions disclosed herein may be embodied in such program instructions, and/or may be implemented in application-specific circuitry (e.g., ASICs or FPGAs) of the computer system. Where the computer system includes multiple computing devices, these devices may, but need not, be co-located. The results of the disclosed methods and tasks may be persistently stored by transforming physical storage devices, such as solid state memory chips and/or magnetic disks, into a different state. In some embodiments, the computer system may be a cloud-based computing system whose processing resources are shared by multiple distinct business entities or other users.

Depending on the embodiment, certain acts, events, or functions of any of the processes or algorithms described herein can be performed in a different sequence, can be added, merged, or left out altogether (e.g., not all described operations or events are necessary for the practice of the algorithm). Moreover, in certain embodiments, operations or events can be performed concurrently, e.g., through multi-threaded processing, interrupt processing, or multiple processors or processor cores or on other parallel architectures, rather than sequentially.

The various illustrative logical blocks, modules, routines, and algorithm steps described in connection with the embodiments disclosed herein can be implemented as electronic hardware (e.g., ASICs or FPGA devices), computer software

15

that runs on general purpose computer hardware, or combinations of both. To clearly illustrate this interchangeability of hardware and software, various illustrative components, blocks, modules, and steps have been described above generally in terms of their functionality. Whether such functionality is implemented as specialized hardware versus software running on general-purpose hardware depends upon the particular application and design constraints imposed on the overall system. The described functionality can be implemented in varying ways for each particular application, but such implementation decisions should not be interpreted as causing a departure from the scope of the disclosure.

Moreover, the various illustrative logical blocks and modules described in connection with the embodiments disclosed herein can be implemented or performed by a machine, such as a general purpose processor device, a digital signal processor (DSP), an application specific integrated circuit (ASIC), a field programmable gate array (FPGA) or other programmable logic device, discrete gate or transistor logic, discrete hardware components, or any combination thereof designed to perform the functions described herein. A general purpose processor device can be a microprocessor, but in the alternative, the processor device can be a controller, microcontroller, or state machine, combinations of the same, or the like. A processor device can include electrical circuitry configured to process computer-executable instructions. In another embodiment, a processor device includes an FPGA or other programmable device that performs logic operations without processing computer-executable instructions. A processor device can also be implemented as a combination of computing devices, e.g., a combination of a DSP and a microprocessor, a plurality of microprocessors, one or more microprocessors in conjunction with a DSP core, or any other such configuration. Although described herein primarily with respect to digital technology, a processor device may also include primarily analog components. For example, some or all of the rendering techniques described herein may be implemented in analog circuitry or mixed analog and digital circuitry. A computing environment can include any type of computer system, including, but not limited to, a computer system based on a microprocessor, a mainframe computer, a digital signal processor, a portable computing device, a device controller, or a computational engine within an appliance, to name a few.

The elements of a method, process, routine, or algorithm described in connection with the embodiments disclosed herein can be embodied directly in hardware, in a software module executed by a processor device, or in a combination of the two. A software module can reside in RAM memory, flash memory, ROM memory, EPROM memory, EEPROM memory, registers, hard disk, a removable disk, a CD-ROM, or any other form of a non-transitory computer-readable storage medium. An exemplary storage medium can be coupled to the processor device such that the processor device can read information from, and write information to, the storage medium. In the alternative, the storage medium can be integral to the processor device. The processor device and the storage medium can reside in an ASIC. The ASIC can reside in a user terminal. In the alternative, the processor device and the storage medium can reside as discrete components in a user terminal.

Conditional language used herein, such as, among others, “can,” “could,” “might,” “may,” “e.g.,” and the like, unless specifically stated otherwise, or otherwise understood within the context as used, is generally intended to convey that certain embodiments include, while other embodiments do not include, certain features, elements and/or steps. Thus,

16

such conditional language is not generally intended to imply that features, elements and/or steps are in any way required for one or more embodiments or that one or more embodiments necessarily include logic for deciding, with or without other input or prompting, whether these features, elements and/or steps are included or are to be performed in any particular embodiment. The terms “comprising,” “including,” “having,” and the like are synonymous and are used inclusively, in an open-ended fashion, and do not exclude additional elements, features, acts, operations, and so forth. Also, the term “or” is used in its inclusive sense (and not in its exclusive sense) so that when used, for example, to connect a list of elements, the term “or” means one, some, or all of the elements in the list.

Disjunctive language such as the phrase “at least one of X, Y, Z,” unless specifically stated otherwise, is otherwise understood with the context as used in general to present that an item, term, etc., may be either X, Y, or Z, or any combination thereof (e.g., X, Y, and/or Z). Thus, such disjunctive language is not generally intended to, and should not, imply that certain embodiments require at least one of X, at least one of Y, or at least one of Z to each be present.

While the above detailed description has shown, described, and pointed out novel features as applied to various embodiments, it can be understood that various omissions, substitutions, and changes in the form and details of the devices or algorithms illustrated can be made without departing from the spirit of the disclosure. As can be recognized, certain embodiments described herein can be embodied within a form that does not provide all of the features and benefits set forth herein, as some features can be used or practiced separately from others. The scope of certain embodiments disclosed herein is indicated by the appended claims rather than by the foregoing description. All changes which come within the meaning and range of equivalency of the claims are to be embraced within their scope.

What is claimed is:

1. A system comprising:

a plurality of computing systems, wherein each computing system hosts a different directory in a plurality of directories; and

a computing resource service provider system comprising one or more computing devices, the computing resource service provider system in communication with the plurality of computing systems, the one or more computing devices configured to execute instructions that cause the computing resource service provider system to:

receive, from a user device, a credential, a client identification, and a directory identifier, wherein the directory identifier identifies a first directory in the plurality of directories, and wherein the client identification is associated with a service that is not associated with the first directory;

determine, by a first computing agent associated with the first directory, whether the credential can be authenticated; and

in response to a determination that the credential can be authenticated,

generate an authentication code,  
transmit the authentication code to the user device,  
instruct the user device to access the service associated with the client identification,

receive the authentication code from the service associated with the client identification,

generate, in response to receiving the authentication code from the service associated with the client

17

identification, an access token based on the authentication code and the directory identifier, generate, in response to receiving the authentication code from the service associated with the client identification, a refresh token, map the access token to the first directory and store the mapping in a database, and transmit the access token and the refresh token to the service associated with the client identification, wherein the access token allows the user device to manage the first directory via the service associated with the client identification.

2. The system of claim 1, wherein the first computing agent is configured to determine whether the credential can be authenticated by instructing a managed directory service to access the first directory.

3. The system of claim 1, wherein the authentication code is valid for a finite period of time, and wherein the one or more computing devices are further configured to execute instructions that cause the computing resource service provider system to generate the access token in response to receiving the authentication code from the service associated with the client identification before the authentication code expires.

4. The system of claim 1, wherein the authentication code is an OAuth code.

5. A computer-implemented method of authenticating an administrator for the management of one or more directories, the method comprising:

as implemented by a computer resource service provider system comprising one or more computing devices, the computer resource service provider system in communication with a plurality of directories, the computer resource service provider system configured with specific executable instructions,

receiving, from a user device, a credential and a directory identifier, wherein the directory identifier identifies a first directory in the plurality of directories, and wherein the user device accessed a first service that is not associated with the first directory;

determining whether the credential can be authenticated; generating, in response to a determination that the credential can be authenticated, an authentication code;

transmitting the authentication code to the user device; instructing the user device to access the first service; receiving the authentication code from the first service accessed by the user device;

generating, in response to receiving the authentication code from the first service, an access token based on the authentication code and the directory identifier;

generating, in response to receiving the authentication code from the first service, a refresh token;

mapping the access token to the first directory and storing the mapping in a database; and

transmitting the access token and the refresh token to the first service accessed by the user device, wherein the access token allows the user device to manage the first directory via the first service accessed by the user device.

6. The computer-implemented method of claim 5, wherein determining whether the credential can be authenticated fur-

18

ther comprises determining whether the credential can be authenticated by instructing a managed directory service to access the first directory.

7. The computer-implemented method of claim 5, wherein the authentication code is valid for a finite period of time.

8. The computer-implemented method of claim 7, wherein generating an access token further comprises generating the access token in response to receiving the authentication code from the service accessed by the user device before the authentication code expires.

9. The computer-implemented method of claim 5, wherein the authentication code is an OAuth code.

10. The computer-implemented method of claim 5, wherein the credential comprises a username and a password.

11. A non-transitory computer storage system comprising a non-transitory storage device, said computer storage system having stored thereon executable program instructions that direct a computer system to at least:

receive, from a user device, a credential and a directory identifier, wherein the directory identifier identifies a first directory in a plurality of directories, and wherein the user device is not associated with any domain of the first directory;

determine whether the credential can be authenticated; generate, in response to a determination that the credential can be authenticated, an authentication code;

transmit the authentication code to the user device;

instruct the user device to access a first service;

receive the authentication code from the first service accessed by the user device in response to a request for an access token to manage the first directory;

generate, in response to receiving the authentication code from the first service, the access token based on the authentication code and the directory identifier;

generate, in response to receiving the authentication code from the first service, a refresh token;

map the access token to the first directory and store the mapping in a database; and

transmit the access token and the refresh token to the first service accessed by the user device, wherein the access token allows the user device to manage the first directory via the first service.

12. The non-transitory computer storage system of claim 11, wherein the executable program instructions further direct the computer system to at least determine whether the credential can be authenticated by accessing the first directory.

13. The non-transitory computer storage system of claim 11, wherein the authentication code is valid for a finite period of time, and wherein the executable program instructions further direct the computer system to at least generate the access token in response to receiving the authentication code from the first service accessed by the user device before the authentication code expires.

\* \* \* \* \*